Generalized Machine Activation Problem

Jian Li, Samir Khuller

University of Maryland, College Park

Jan. 2011

Problem Definition

- Unrelated Machine Scheduling:
  - M: the set of machines
  - J: the set of jobs
  - p_{ij}: processing time of job j on machine i
  - Goal: find an assignment s.t. the makespan is minimized

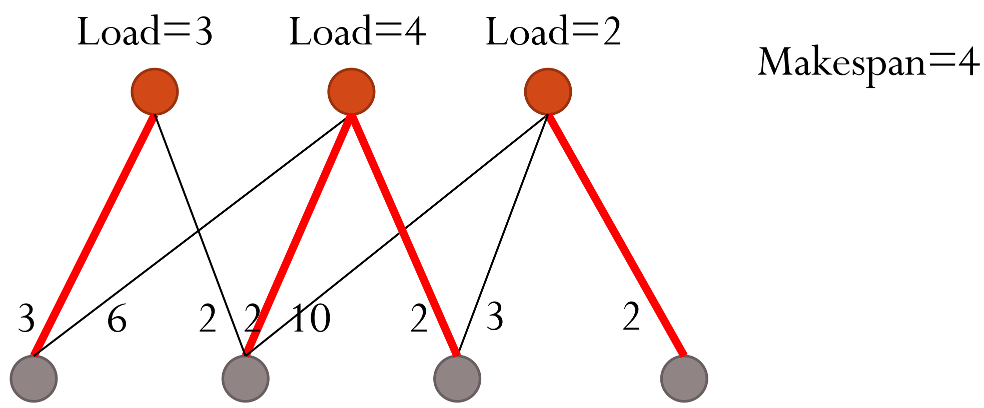

Problem Definition

- **Generalized Machine Activation (GMA):**
  - Machine Activation Cost:

    $w_i(x)$: activation cost function of machine i

    --- A function of the load of machine i

    --- Non-decreasing and piecewise linear

    --- Left-Continuous

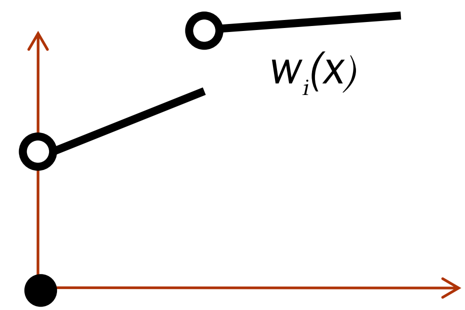

  - Assignment Cost

    a_{ij}: the cost of assigning job j to machine i

  - Objective

    Find an assignment such that the total cost (i.e., machine activation cost plus assignment cost) is minimized

Problem Definition

- GMA generalizes …
  - Machine Activation Problem [Khuller,Li,Saha'10]
    - The activation cost for each machine is fixed; We require the makespan is at most *T* and minimize the total cost
    - $w_i(x)=w_i$ for $0<x<=T$, and $w_i(x)=\infty$ for $x>T$
  - Universal Facility Location [Hajiaghayi,Mahdian,Mirrokni '99] [Mahdian, Pal '03]
    - $p_{ij}=1$ for all i,j, i.e., the activation cost (i.e., facility opening cost) of machine *i* is an increasing function of the number of jobs assigned to *i*
  - Generalized Submodular Covering [Bar-Ilan,Kortsarz,Peleg'01]
    - GSC generalizes the average cost center problem, the fault tolerant facility location problem and the capacitated facility location problem.

Our Results

THM: There is a polynomial time algorithm that finds a fractional assignment such that n-ε jobs are (fractionally) satisfied and the cost is at most $\ln(n/\varepsilon)+1$ times the optimal solution.

- Machine Activation Problem
  - Bicriteria approximation: (makespan, total cost)
  - Previous results:
    - $(2+\varepsilon, 2\ln(2n/\varepsilon)+5)$ [Fleischer'10], $(3+\varepsilon, (1/\varepsilon)\ln(n)+1)$ [KLS'10]
    - No assignment cost: $(2+\varepsilon, \ln(n/\varepsilon)+1)$ [Fleischer'10], $(2, \ln(n)+1)$ [KLS'10]
  - Our results
    - $(2, (1+o(1))\ln(n))$

Our Results

- Universal Facility Location
  - Previous results:
    - Metric: Constant approximations [Mahdian, Pal ’03] [Vygen ’07]
    - Non-metric: Open [Hajiaghayi,Mahdian,Mirrokni ’99] [Mahdian, Pal ’03]
  - Our results
    - Non-metric: *(ln(n)+1)*-approximation
- Generalized Submodular Covering
  - Previous results:
    - *O(ln nM)*-approximation where *M* is the largest integer in the instance
  - Our results:
    - *ln(D)*-approximation where *D* is the total demand

Our Results

Machine Activation with Linear Constraints

- Each machine has a fixed activation cost
- For each machine, the set of jobs assigned to it must satisfy a set of *d* linear constraints

$$\sum_{j \in J} p_{ijk} x_{ij} \cdot \quad T_{ik} \qquad i \in M, k = 1, 2, ..., d$$

E.g., makespan constraint, degree constraint …

- **THM:** For any *ε>0*, there is a poly-time algorithm that returns an integral schedule X,Y such that
  1. (1) $\sum_{j \in J} p_{ijk} X_{ij} \cdot \quad (2d + \epsilon) T_{ik}$ for each i and $1 \cdot \quad k \cdot \quad d$;
  2. (2) $\mathbb{E}[\sum_{i \in M} \omega_i Y_i] \cdot \quad O(\frac{1}{\epsilon} \log n) \sum_{i \in M} \omega_i y_i$.
- This matches the previous bound for *d=1* [KLS10]

Outline

- Greedy for Universal Facility Location
- Greedy for Generalized Machine Activation
- Final Remarks

Greedy for UFL

- A set of facilities (machines) and clients (jobs)
- Facility opening cost $w_i(u_i)$ which is a non-decreasing function of the load of facility *i* (load= #clients assigned to it)
- Assignment cost: a_{ij}

- $\boldsymbol{u}$: the load vector
- $\pi(\boldsymbol{u})$: min. assignment cost under load vector $\boldsymbol{u}$
- $C(\boldsymbol{u})=\sum_i w_i(u_i) + \pi(\boldsymbol{u})$

  -- $\pi(u)$ can be computed via a min-cost flow

$\boldsymbol{u}=<0,1,2,0>$

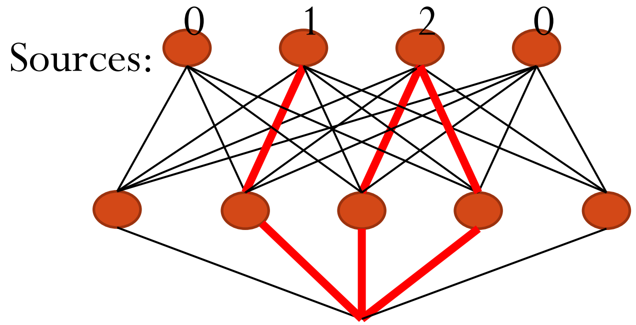

Greedy for UFL

- $\boldsymbol{u}$: the load vector
- $C(\boldsymbol{u})=\sum_i w_i(u_i) + \pi(\boldsymbol{u})$
- $\mathbf{e}_i$= <0,…,1,…,0>

  (The ith entry)

- **GREEDY-UFL**

  Repeat

  -- choose the machine i and integer $k>0$ such that

$$\rho(\mathbf{u}, i, k) = \frac{C(\mathbf{u}+k\mathbf{e}_i)-C(\mathbf{u})}{k}$$

  is minimized.

  Until all jobs are served (i.e., $|\boldsymbol{u}|=n$)

Greedy for UFL

- **Analysis:**
  - We would like to show

$$\min_{i,k} \rho(\mathbf{u}, i, k) \cdot \frac{C(\mathbf{u}^*)}{n - |\mathbf{u}|}$$

where **u*** is the optimal load vector

Lemma: For any load vector $\mathbf{u}$, there exists $\widetilde{\mathbf{u}}$ such that

1. $\mathbf{u} \cdot \widetilde{\mathbf{u}} \cdot \max(\mathbf{u}, \mathbf{u}^*)$
2. $\pi(\widetilde{\mathbf{u}}) \cdot \pi(\mathbf{u}^*) + \pi(\mathbf{u})$
3. $|\widetilde{\mathbf{u}}| = n$

$\mathbf{u} = \langle 0, 0, 1, 3, 3 \rangle$

$\mathbf{u}^* = \langle 4, 4, 2, 2, 0 \rangle$

$\widetilde{\mathbf{u}} = \langle 1, 2, 2, 3, 3 \rangle$

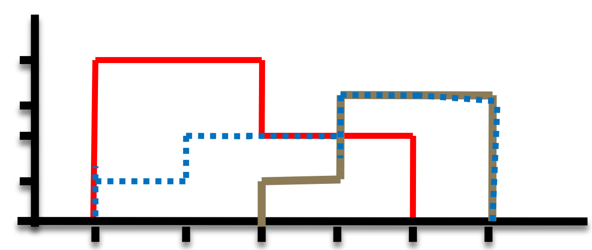

Greedy for UFL

- **Analysis Cont:**

f (or $\widetilde{f}$) is the optimal flow corresponding to $\mathbf{u}$ (or $\widetilde{\mathbf{u}}$)

Consider the flow $g = \widetilde{f} - f$

(1) We can easily show *g* is a feasible flow in the residual graph w.r.t. *f*

(2) Apply the conformal path decomposition to *g*.

(3) Divide the paths into groups (g_1, g_2,…) base on the sources of the paths (indicated by colors)

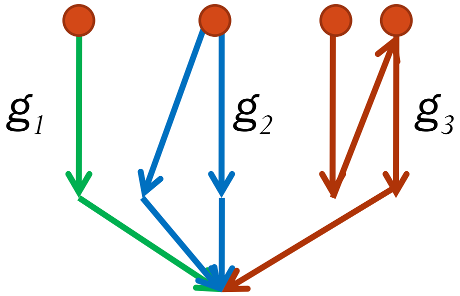

Such a structure is due to the fact that $\widetilde{\mathbf{u}} \geq \mathbf{u}$

Greedy for UFL

- **Analysis cont.**

1. $\mathbf{u} \cdot \ \widetilde{\mathbf{u}} \cdot \ \max(\mathbf{u}, \mathbf{u}^*)$
2. $\pi(\widetilde{\mathbf{u}}) \cdot \ \pi(\mathbf{u}^*) + \pi(\mathbf{u})$

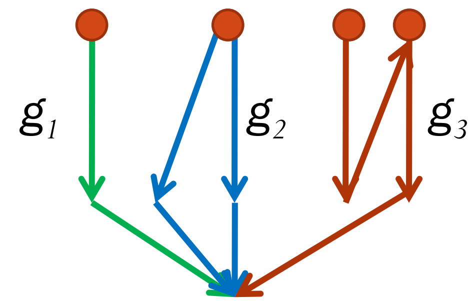

$$\sum_i c(g_i) = c(g) = c(\widetilde{f}) - c(f) = \underbrace{\pi(\widetilde{\mathbf{u}}) - \pi(\mathbf{u}) \cdot \ \pi(\mathbf{u}^*)}_{\text{Lemma (2)}}$$

Therefore,

$$\sum_i \left(c(g_i) + \underbrace{w(\tilde{u}_i) - w(u_i)) \leq \sum_i w(u_i^*)}_{\text{Lemma (1)}} + \pi(\mathbf{u}^*) \leq C(\mathbf{u}^*)\right.$$

Greedy for UFL

- **Analysis cont.**

1. $\mathbf{u} \cdot \tilde{\mathbf{u}} \cdot \max(\mathbf{u}, \mathbf{u}^*)$
2. $\pi(\tilde{\mathbf{u}}) \cdot \pi(\mathbf{u}^*) + \pi(\mathbf{u})$

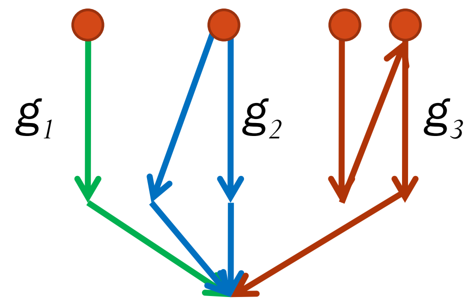

$$\min_{i,k} \rho(\mathbf{u}, i, k) \cdot \min_i \frac{c(g_i)+w(\tilde{u})-w(u_i)}{r(g_i)} \cdot \frac{C(\mathbf{u}^*)}{n-|\mathbf{u}|}$$

g_i is feasible on the residual graph w.r.t. f

Greedy for UFL

1. $\mathbf{u} \cdot \ \widetilde{\mathbf{u}} \cdot \ \max(\mathbf{u}, \mathbf{u}^*)$
2. $\pi(\widetilde{\mathbf{u}}) \cdot \ \pi(\mathbf{u}^*) + \pi(\mathbf{u})$

Pf of the lemma (sketch):

- f (or f^*) is the optimal flow corresponding to $\mathbf{u}$ (or $\mathbf{u}^*$)
  Consider the flo $g = f^* - f$
- Divide the paths into two groups g_1 and g_2 (indicated by colors)
- Consider flow $\widetilde{f} = f + g_1$

  Only need to show $c(g_1) <= c(f^*)$
  Notice that $f^* - g_1 = f + g_2$, which is a feasible flow on the original graph

$\mathbf{u} = \langle 0, 0, 1, 2 \rangle$

$\mathbf{u}^* = \langle 2, 2, 2, 0 \rangle$

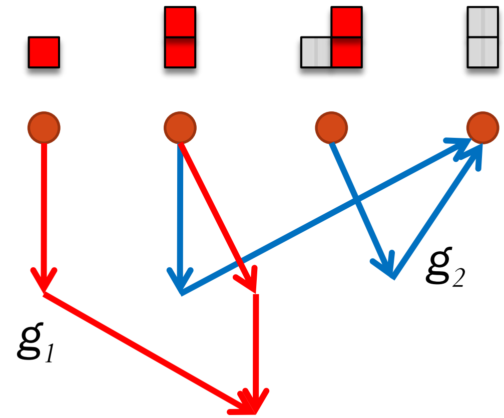

Outline

- Greedy for Universal Facility Location
- Greedy for Generalized Machine Activation
- Final Remarks

Algorithm for GMA

- The algorithm is similar to GREEDY-UFL, except that
  - The optimal (fractional) assignment cost can be computed via a *generalized flow* computation

If *1* unit of flow goes in, γ_e units of flow go out

  - The flow augmented in each iteration is not necessarily integral anymore. Therefore, we need to put a lower bound on it to ensure polynomial running time.
  - Finding the optimal ratio can be formulated as a *linear-fractional program*

Algorithm for GMA

- Conformal decomposition for generalized flows: a generalized flow can be decomposed into bi-cycles.

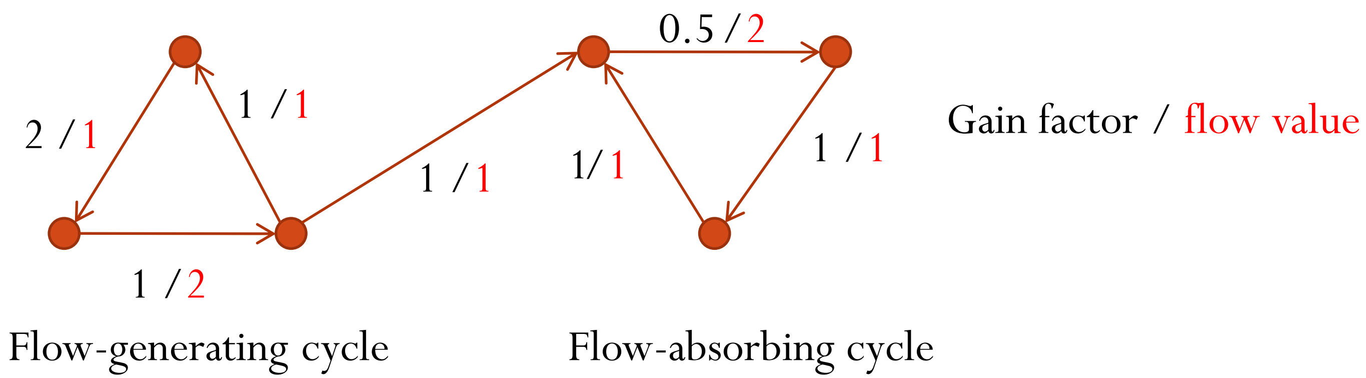

- A **cleanup** procedure to eliminate negative bi-cycles without increasing the total cost (for technical reasons)

Final Remarks

- We give two proofs of the supermodularity of the generalized flow (first proved in [Fleischer'10]).
  - The first one is based on the conformal decomposition of a generalized flow
  - The second one is based on the conformal decomposition of the dual LP solution (which is not a flow)

- How to handle non-increasing machine activation cost?
  - Lower-bounded facility location [Karger, Minkoff '00][Guha, Meyerson, Munagala'00][Svitkina'08]

Thanks

Texpoint 3.2.1

- SODA 2011
- 22-23 min talk (25 min slot)

$\mathbf{u} = \langle 0, 1, 2, 0 \rangle$

$\widetilde{\mathbf{u}} = \langle 1, 2, 3, 0 \rangle$

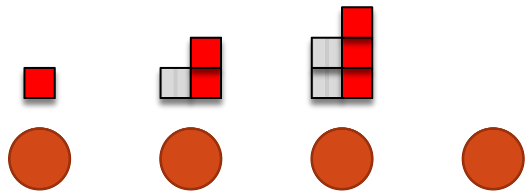

1. $\mathbf{u} \cdot \widetilde{\mathbf{u}} \cdot \max(\mathbf{u}, \mathbf{u}^*)$
2. $\pi(\widetilde{\mathbf{u}}) \cdot \pi(\mathbf{u}^*) + \pi(\mathbf{u})$

Greedy for Set Cover

- Set Cover:
  - A set *U* of elements
  - A family of subsets of *U*, each associated with a weight
  - Goal: find a min-weight covering of *U*

- GREEDY-SC

  Repeat

  -- choose the set *s* minimizing $\rho(s) = \frac{w(s)}{|s \cap U_i|}$

  -- $U_{i+1} = U_i - S$

  -- *i=i+1*

  Until U_i is empty

THM: GREEDY-SC is an *ln(n)*-approximation.

Greedy for Set Cover

- Analysis: Suppose we choose s_i at step *i*

  We would like to show

$$\rho(s_i) = \frac{w(s)}{|s \cap U_i|} \cdot \frac{OPT}{n - |U_i|}$$

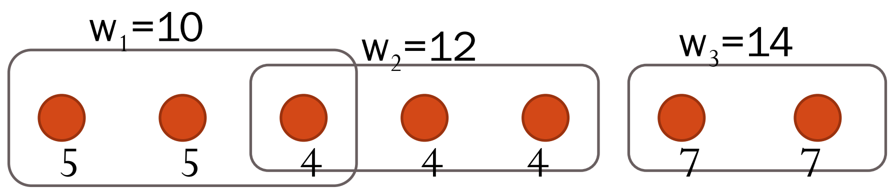

Then we have that our cost is

$$\sum_i \rho(s_i)|s_i \cap U_i| \cdot OPT \sum_i \frac{1}{n - |U_i|} \cdot OPT \sum_{i=1}^{n} \frac{1}{i} \cdot \ln n OPT$$